Argonne NATIONAL LABORATORY

ALCF Service Offerings

David E. Martin
Manager, User Services and Outreach
dem@alcf.anl.gov

ALCF Getting Started Workshop
January 18, 2011

U.S. DEPARTMENT OF ENERGY

Argonne Leadership Computing Facility

- ALCF was established in 2006 at Argonne to provide the computational science community with a leading-edge computing capability dedicated to breakthrough science and engineering
- One of two DOE national Leadership Computing Facilities (the other is the National Center for Computational Sciences at Oak Ridge National Laboratory)
- Supports the primary mission of DOE's Office of Science Advanced Scientific Computing Research (ASCR) program to discover, develop, and deploy the computational and networking tools that enable researchers in the scientific disciplines to analyze, model, simulate, and predict complex phenomena important to DOE.

ALCF Compute and Storage Resources

- *Intrepid* - ALCF Blue Gene/P System:
  - 40,960 nodes / 163,840 PPC cores
  - 80 Terabytes of memory
  - Peak flop rate: 557 Teraflops
  - Linpack flop rate: 450.3
  - #13 on the Top500 list
- *Eureka* - ALCF Visualization System:
  - 100 nodes / 800 2.0 GHz Xeon cores
  - 3.2 Terabytes of memory
  - 200 NVIDIA FX5600 GPUs
  - Peak flop rate: 100 Teraflops
- Storage:
  - 6+ Petabytes of disk storage with an I/O rate of 80 GB/s (GPFS and PVFS)
  - 5+ Petabytes of archival storage, 10,000 volume tape archive (HPSS)

ALCF Compute Resources

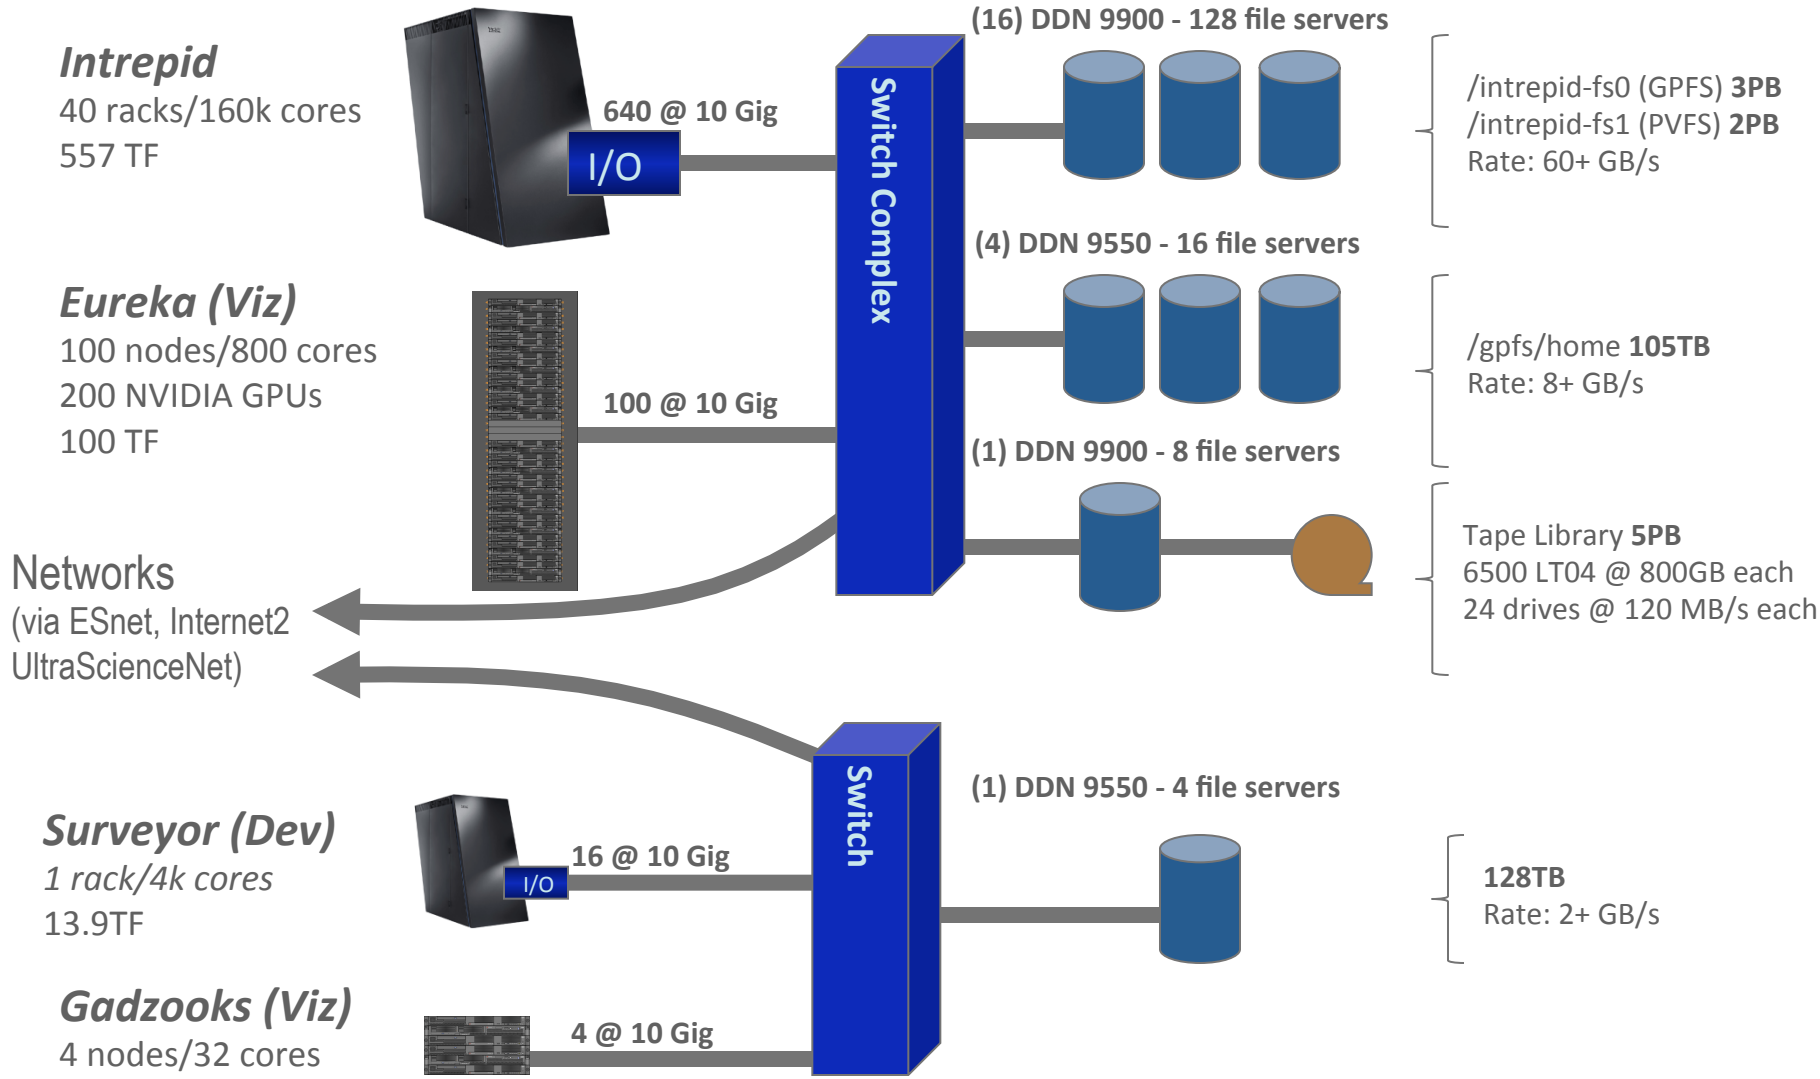

Percentage of compute hours used by scientific discipline
January–October 2010
Total hours used: 849 Million

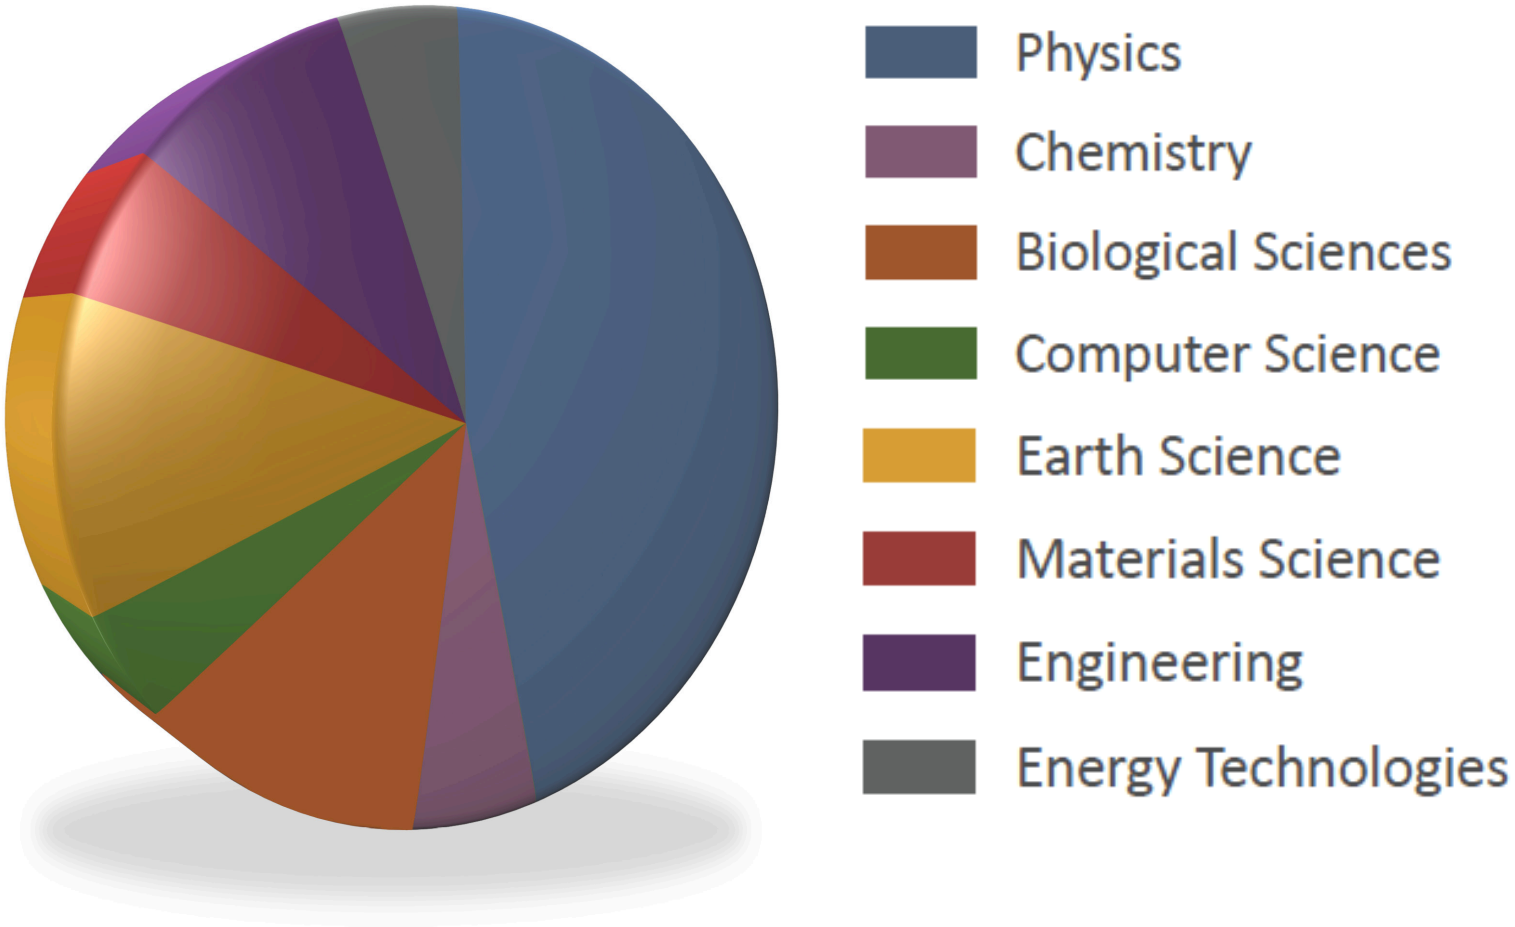

Variety of Allocation Programs

60%	30%	10%	
Innovative and Novel Computational Impact on Theory and Experiment (INCITE)	**ASCR Leadership Computing Challenge Program (ALCC)**	**Early Science Program (ESP)**	**Discretionary Projects**
ALCF resources are available to researchers as part of the U.S. Department of Energy's INCITE program. Established in 2003, the program encompasses high-end computing resources at Argonne and other national laboratories. The INCITE program specifically seeks out computationally intensive, large-scale research projects with the potential to significantly advance key areas in science and engineering. The program encourages proposals from universities, other research institutions, and industry. It continues to expand, with current research applications in areas such as chemistry, combustion, astrophysics, genetics, materials science and turbulence.	Open to scientists from the research community in academia and industry, the ALCC program allocates resources to projects with an emphasis on high-risk, high payoff simulations in areas directly related to the Department's energy mission, national emergencies, or for broadening the community of researchers capable of using leadership computing resources. Projects are awarded an ALCC allocation based on a peer review for scientific merit and computational readiness.	Allocations through the Early Science Program (ESP) provide researchers with preproduction hours (between system installation and full production) on the ALCF's next-generation, 10 petaflops IBM Blue Gene system. This early science period provides projects with a significant head start for adapting to the new machine and access to substantial computational time. During this shakedown period, users assist in identifying the root causes of any system instabilities, and work with ALCF staff to help develop solutions. More than four billion core hours are allocated through ESP.	Discretionary allocations are "start up" awards made to potential future INCITE projects. Projects must demonstrate a need for leadership-class resources. Awards may be made year round to industry, academia, laboratories and others, and are usually between three and six months in duration. The size of the award varies based on the application and its readiness/ability to scale; awards are generally from the low tens of thousands to the low millions of hours.

DOE INCITE Program

Innovative and Novel Computational Impact on Theory and Experiment

- **Solicits large computationally intensive research projects**
  - To enable high-impact scientific advances
  - Call for proposal opened once per year (call closed 6/30/2010)
  - INCITE Program web site: www.er.doe.gov/ascr/incite
- **Open to all scientific researchers and organizations**
  - Scientific Discipline Peer Review
  - Computational Readiness Review
- **Provides large computer time & data storage allocations**
  - To a small number of projects for 1-3 years
  - Academic, Federal Lab and Industry, with DOE or other support
- **Primary vehicle for selecting principal science projects for the Leadership Computing Facilities**
  - 60% of time at Leadership Facilities
- **In 2010, 35 INCITE projects allocated more than 600M CPU hours at the ALCF**

DOE ALCC Program

ASCR Leadership Computing Challenge

- Allocations for projects of special interest to DOE with an emphasis on high risk, high payoff simulations in areas of interest to the department's energy mission
  - 30% of the core hours at Leadership Facilities
- Awards granted in June, 2010
  - 10 awards at ALCF in 2010 for 300+ million core hours
- ALCC applications submitted from August 1, 2010 through February 15, 2011 will be considered for allocation in 2011
  - http://www.science.doe.gov/ascr/facilities/alcc.html

Director's Discretionary Allocations

- Time is available for projects without INCITE or ALCC allocations!
- ALCF Discretionary allocations provide time for:
  - Porting, scaling, and tuning applications
  - Benchmarking codes and preparing INCITE proposals
  - Preliminary science runs prior to an INCITE award
  - Early Science Program
- To apply go to the ALCF allocations page
  - www.alcf.anl.gov/support/gettingstarted

ALCF Organizational Structure

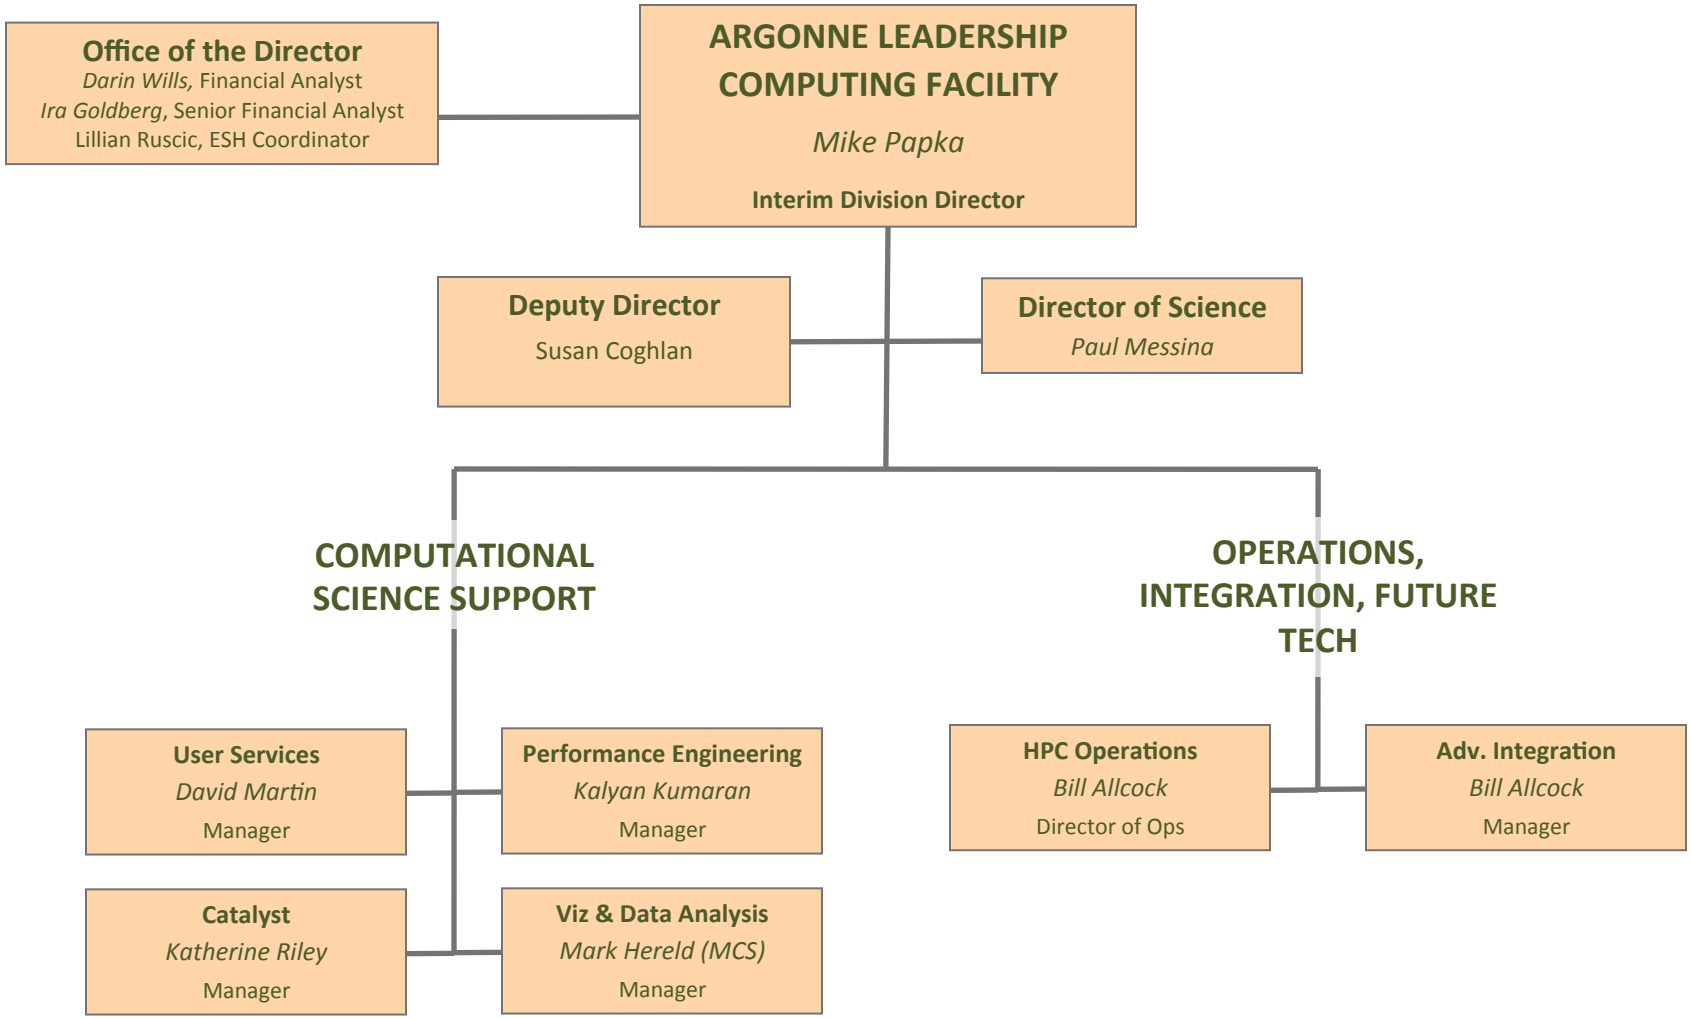

Performance Engineering Team

The mission of the Performance Engineering group is to help ALCF users achieve the best performance out of their applications. To this end, ALCF performance engineers work closely with the users in porting, tuning and parallelizing their applications on ALCF computers. They also assist in resolving performance-inhibiting I/O issues.

ALCF Performance Engineers have extensive experience in:

- Porting, performance tuning and parallelizing of scientific applications and other software
- Computer architectures
- Computational algorithms
- I/O

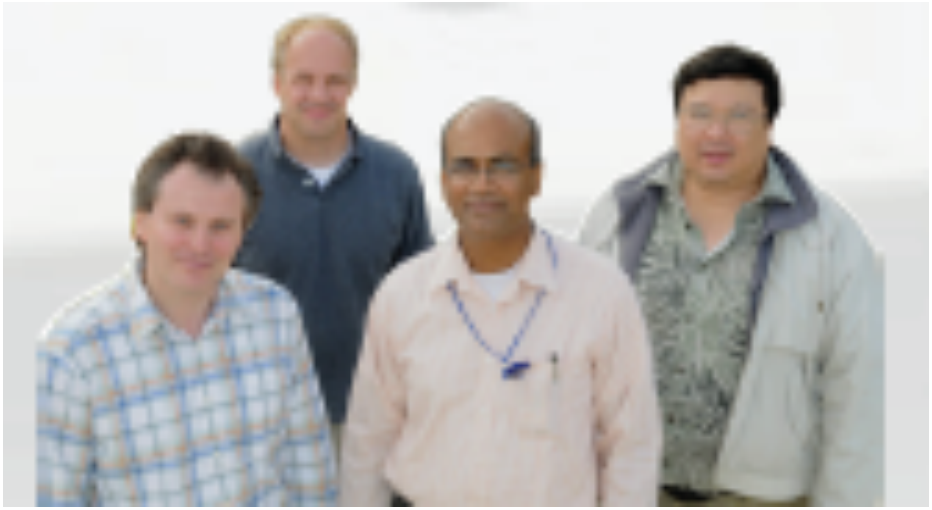

Catalyst Team

The Catalyst team provides key projects with an assigned expert, the "Catalyst," to maximize and accelerate research. Catalysts are computational scientists that are experts in their fields: computational chemistry, physics, fluid dynamics, astrophysics, etc., and actively participate in the research projects.

In addition to diagnosing problems using their knowledge of the scientific calculations being done, Catalysts also provide:

- Assistance with algorithm development and scaling efforts
- A single point of contact for key research initiatives
- Tailored services for unique project requirements
- Full project lifecycle assistance
- Value-added services and support in conjunction with ALCF resources

Data Analytics and Visualization Team

The Data Analytics and Visualization team has expertise in tools and methods for high-performance post processing of large datasets, interactive data exploration, batch visualization, and production visualization.

Team members help users with their visualization and analysis needs using ALCF high-performance resources and a suite of tools maintained for these purposes.

- Production tools for high performance visualization (ParaView, VisIt)
- Analysis tools (R, MATLAB)
- Presentation graphics (PowerPoint, Keynote, Final Cut Pro)
- The ALCF Data Analytics and Visualization team has strong connections to Argonne's Mathematics and Computer Science research and development in the area of visualization and analysis.

Operations Team

The ALCF Operations team consists of the Systems Group and the Advanced Integration Group.

The Systems Group is responsible for:

- Hardware maintenance
- Software maintenance
- Resolution of user tickets related to system issues
- Responding to requests for new software or versions of software
- Developing systems tools, particularly ones related to the unique system architectures and scale of ALCF resources

The Advanced Integration Group is responsible for:

- Ensuring the entire system software stack works together
- Assisting with I/O performance issues
- Bug fixes and feature requests for systems software

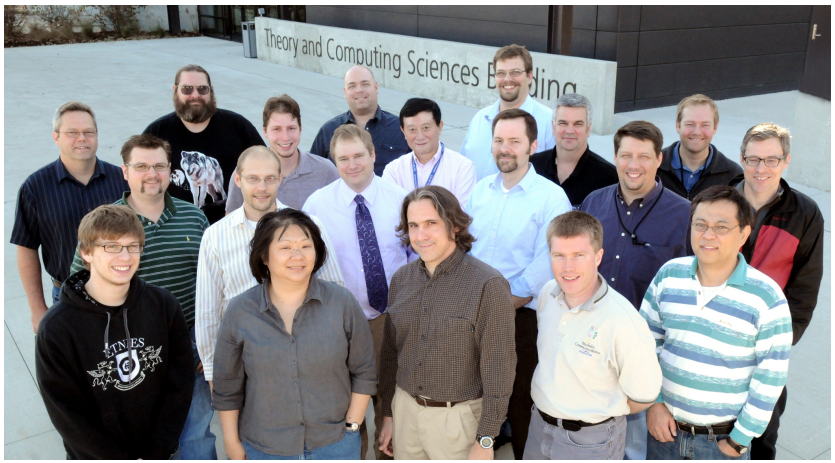

User Services and Outreach

The USO team provides frontline services and support to existing and potential ALCF users. The team also provides marketing and outreach to users, DOE and the broader community.

USO provides the following services:

- Account/Project Administration
  - Account management, foreign national clearance, secure access
- Training and Education
  - Documentations, workshops
- Help Desk
  - Problem resolution, User Communications
- User Outreach
  - Tours, Classes, Conferences
- Support@alcf.anl.gov

Integrated Services Model

- Startup assistance
- User administration assistance
- Job management services
- Technical support

- Science project management
- Assistance with proposals, planning, reporting
- Collaboration within science domains
- ALCF point of coordination

- Workshops and seminars
- Customized training programs
- On-line content and user guides
- Educational and outreach programs
- Reporting and promotion

- Performance engineering
- Application tuning
- Data analytics
- Visualization
- Data management services

ALCF Services - Start

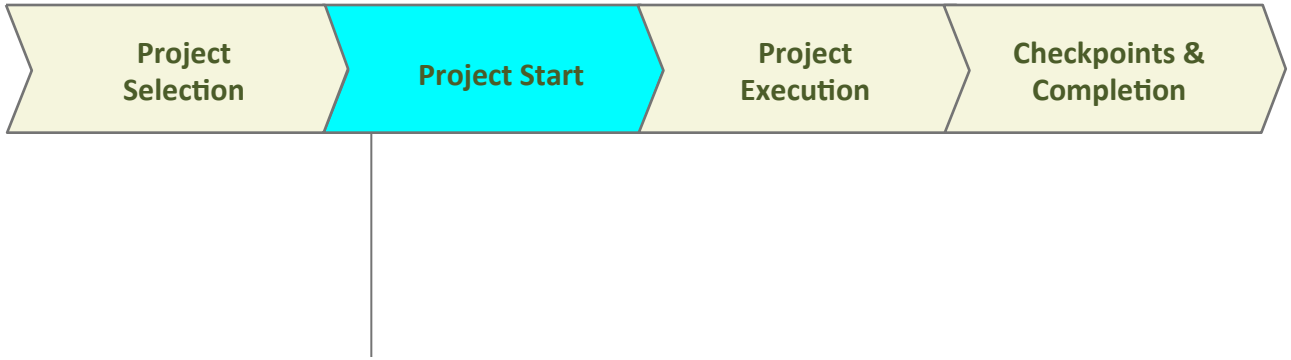

- Create Project and Allocation in ALCF Management System
- Sign Institution-Level User Agreements, Foreign-National Access
- Appoint ALCF Catalyst
- Introduce Research Team to ALCF Team Members
- Review ALCF Service Offerings
- Discuss Areas of Potential Assistance and Collaboration
- “Getting Started” Workshop

ALCF Services - Execution

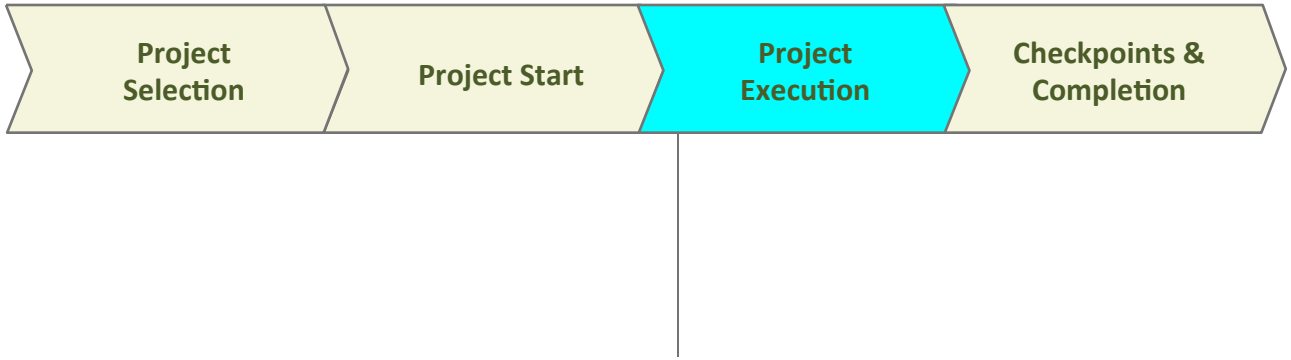

- Collaboration with ALCF Computational Scientists to Develop Computation Strategy
- Collaboration with ALCF Performance Engineering and Data Analytics Specialists to Optimize Algorithms and Code
- Ongoing User Support and Tracking
- Workshops, INCITE User Calls, Training
- Host Visiting Scientists

ALCF Services – Checkpoints and Completion

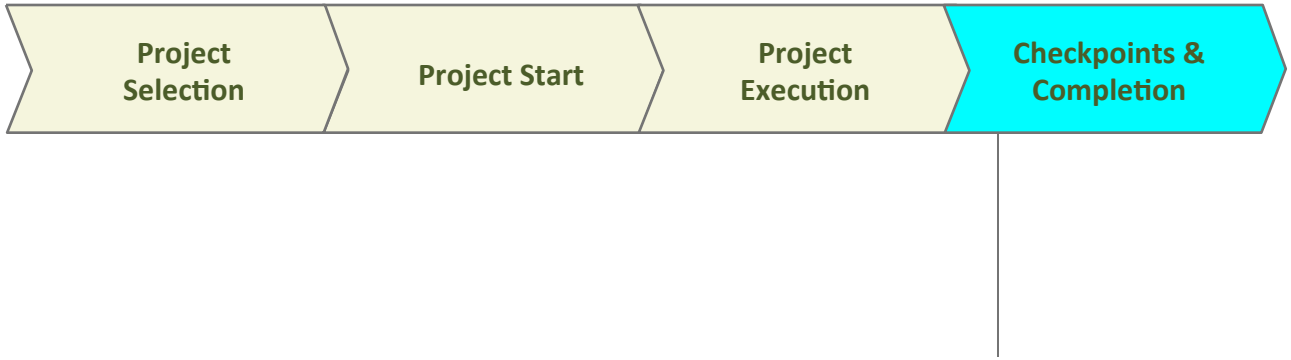

- Gather Usage Statistics and Performance Data
- Monitor Allocation Usage and Planning for Next Award Cycle
- Complete Surveys and Report Significant Milestones
- Generate Regular Project Report for ALCF and DOE
- Highlight Work at DOE and SC Meetings
- Complete Project Close-Out and Archives of Data

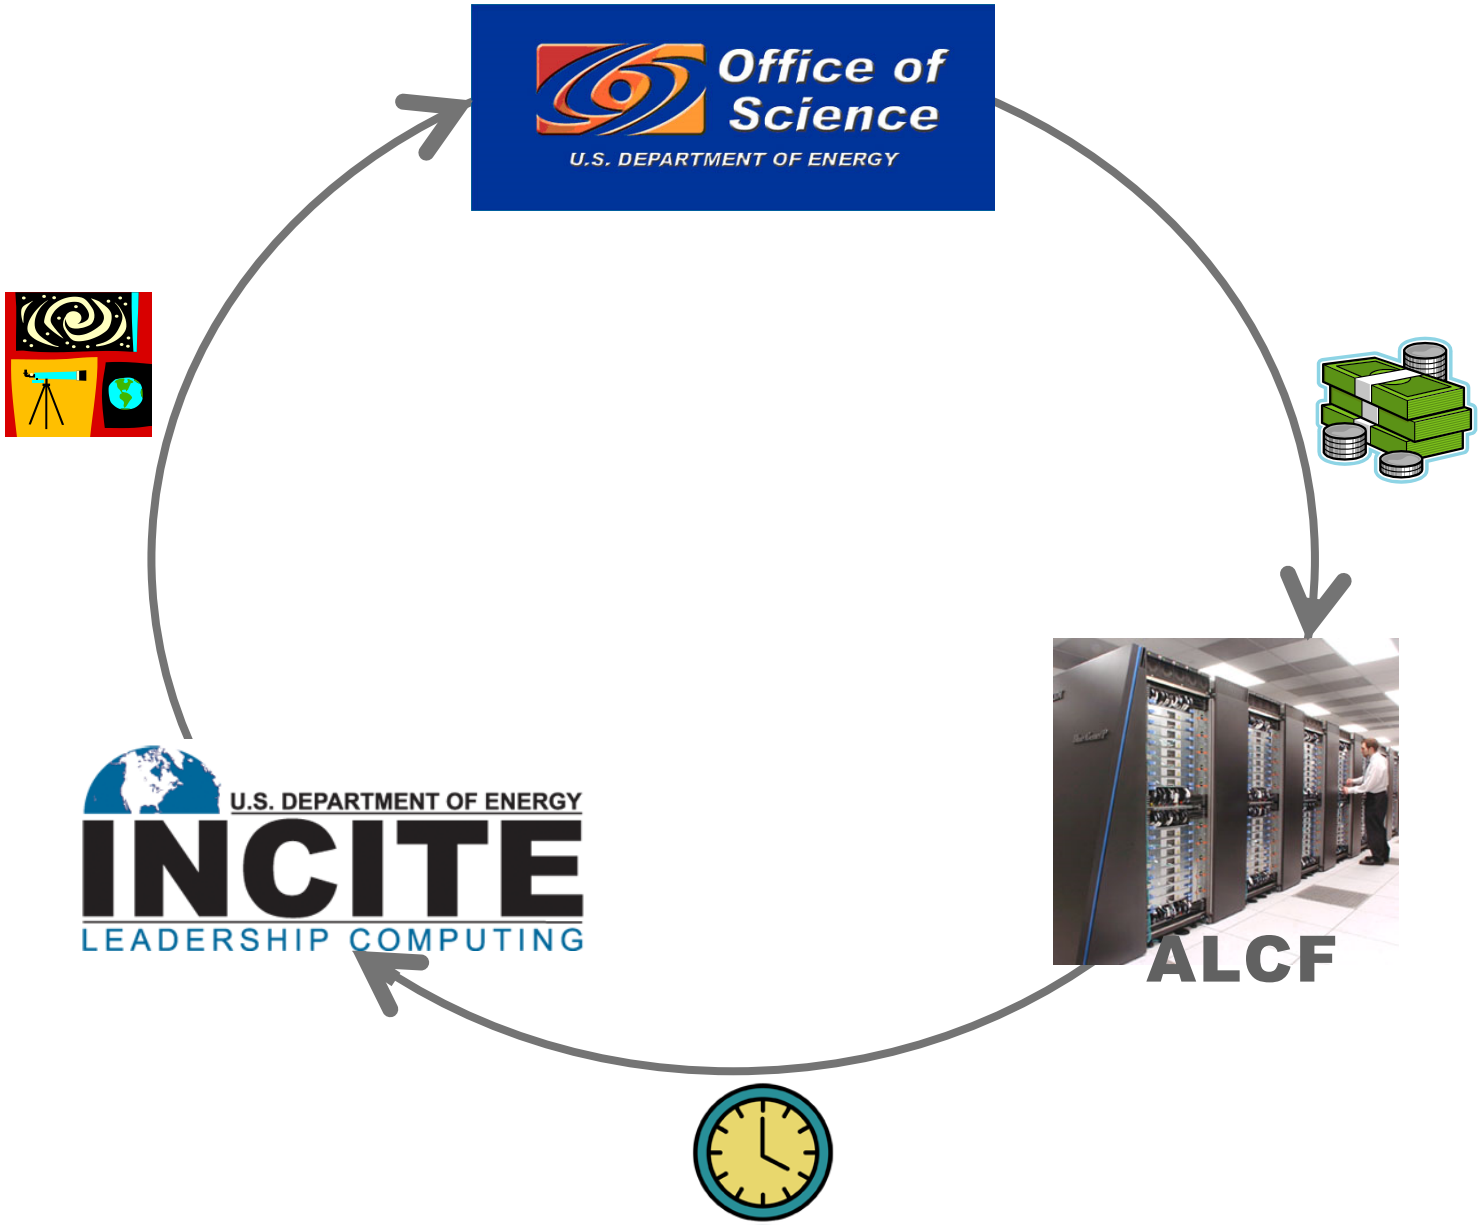
Office of Science
U.S. DEPARTMENT OF ENERGY
U.S. DEPARTMENT OF ENERGY
INCITE
LEADERSHIP COMPUTING
ALCF

Reporting

- Quarterly Updates
  - Coordinated Between OLCF/ALCF
  - Used for Consolidated Report to DOE/ASCR
- Publications
  - Acknowledge DOE/ALCF and Send Us A Copy
- Highlights for DOE
  - Need Graphics, Photos, Images, Videos, Animations, ...
  - For DoE Publications, Web Site, Presentations, SC

Partnership

Projects at ALCF Are A Partnership Between

- DOE: Direction, Funding and Oversight
- INCITE: Resource Allocations and Coordination
- Argonne: Compute Resources and Computational Science
- Science Projects: Domain Expertise and Complex Challenges

Some Ideas for the Future

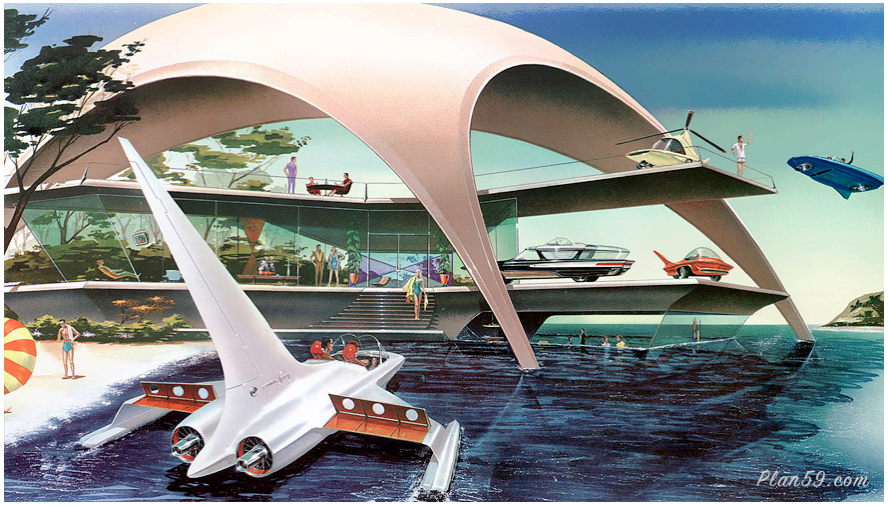

- Better Web Site
  - Better Organization
  - Portals
  - Real-Time Status
- Community
  - Community Forum
  - User Group, User Meeting
  - Community-Based Allocations
- More Specialized Workshops

Questions?